The science behind the report:

Add new-generation Dell EMC PowerEdge MX servers to your VMware Cloud Foundation infrastructure with ease

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report Add new-generation Dell EMC PowerEdge MX servers to your VMware Cloud Foundation infrastructure with ease.

We concluded our hands-on testing on March 7, 2021. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on March 7, 2021 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to http://facts.pt/calculating-and-highlighting-wins. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our testing

Steps to add a cluster of three new-generation Dell EMC™ PowerEdge™ MX servers		Time
1	Insert first MX750c sled into MX7000 chassis, and configure BIOS, power, iDRAC, and disk controller settings	0:04:15
2	Use Dell OpenManage™ Enterprise web interface to create a template from the MX750c sled, and apply VLAN settings for Vmware Cloud Foundation	0:02:50
3	Insert two remaining MX750c sleds. Once online, deploy the template, specifying the ESXi ISO to mount from a CIFS share during the process.	0:25:59
4	Connect to the iDRAC remote consoles for each MX750c sled. Install and configure ESXi, setting management network values, VLANs, enabling command line access, and time services.	0:19:17
5	Commission the MX750c workload hosts in VMware® SDDC Manager. Create a VI Workload domain using the newly commissioned hosts	1:29:00

System configuration information

Table 2: Detailed configuration information on the servers we tested.

System configuration information	3 x Dell EMC PowerEdge MX740c	3 x Dell EMC PowerEdge MX750c
BIOS name and version	Dell 2.9.4	Dell 0.4.2
Non-default BIOS settings	Virtualization performance mode	Virtualization performance mode
Operating system name and version/build number	VMware ESXi 7.0.1 17551050 U1 P30	VMware ESXi 7.0.1 17551050 U1 P30
Date of last OS updates/patches applied	03/07/21	03/07/21
Power management policy	Performance	Performance
Processor		
Number of processors	2	2
Vendor and model	Intel® Xeon® Gold 6230	Intel Xeon Gold 6330
Core count (per processor)	20	28
Core frequency (GHz)	2.10	2.00
Stepping	7	6
Memory module(s)		
Total memory in system (GB)	192	512
Number of memory modules	12	16
Vendor and model	Hynix HMA82GR7AFR8N-VK	Hynix HMAA4GR7AJR8N-XN
Size (GB)	16	32
Type	PC4-21300	PC4-23400
Speed (MHz)	2,666	2,933
Speed running in the server (MHz)	2,666	2,933
Storage controller		
Vendor and model	VMware NVMe™ PCIe®	VMware NVMe PCIe
Driver version	1.2.3.9-2vmw.701.0.0.16850804	1.2.3.9-2vmw.701.0.0.16850804
Local storage (type A)		
Number of drives	4	4
Drive vendor and model	Dell Ent NVMe AGN MU U.2 1.6TB	Dell Ent NVMe AGN MU U.2 3.2TB
Drive size (GB)	1,600	3,200
Drive information (speed, interface, type)	U.2 NVMe, 8GT/s	U.2 NVMe, 16GT/s
Network adapter		
Vendor and model	Intel Ethernet 25G 2P XXV710 Mezz	Broadcom Adv Quad 25Gb Ethernet
Number and type of ports	2 x 25GbE	4 x 25GbE
Driver version	1.8.1.123-1vmw.701.0.0.16850804	216.0.50.0-16vmw.701.0.0.16850804

Table 3: Configuration information for the server enclosure we tested.

System configuration information	Dell EMC MX7000 Modular Chassis
Number of management modules	2
Management module firmware revision	1.30.00
I/O modules	
Vendor and model number	Dell EMC Networking MX9116N Fabric Switching Engine
I/O module firmware revision	10.5.1.7.273
Number of modules	2
Occupied bay(s)	A1, A2
Power supplies	
Vendor and model number	Dell 0H7TFGA02
Number of power supplies	6
Wattage of each (W)	3,000
Cooling fans	
Vendor and model number	Dell 0FHH0KA00
Number of fans	9

Table 4: Configuration information for the network switches we tested.

System configuration information	2 x Dell EMC Networking S4048-ON
Firmware revision	10.5.1.4.249
Number and type of ports	48 x SFP+ 10GbE, 6 x QSFP+ 40GbE
Number and type of ports used in test	12 x SFP+ 10GbE, 8 x QSFP+ 40GbE

How we tested

MX-series compute sled deployment in VMware Cloud Foundation

Configuring MX7000 with MX740c and MX750c Compute Sleds

This configuration assumes a pre-existing, functional management domain in VMware Cloud Foundation networked to the MX7000 chassis via MX9116N I/O modules, using redundant VLAN uplink trunk connections between the management cluster and the MX chassis by way of S4048-ON switches in VLT peer configuration.

1. Insert first compute sled into MX7000 chassis and power on, wait for the chassis to acknowledge sled presence in OpenManage Enterprise Modular web interface.
2. Wait for Health Check to complete and the iDRAC to come online.
3. Connect to the iDRAC and configure BIOS, power, iDRAC, and disk controller settings as desired for configuration.
4. In OME, select Configuration→Templates.
5. Select Create Template→From Reference Device.
6. Name the template, and select Elements to Clone: iDRAC, BIOS, System, NIC, Lifecycle Controller.
7. Click Select Device, then select the new sled added to the chassis.
8. Click finish twice, and wait for template creation to complete.
9. Select the new template, and click Edit Network.
10. Click VLANs in the left column.
11. For each NIC, set the VLAN for the Untagged Network and select the VLANs for the tagged networks, including all required VMware Cloud Foundation VLANs, click Finish.
12. Insert the remaining compute sleds into the chassis and wait for them to come online.
13. Wait for the health checks to complete and for the iDRACs to come online.
14. Select the previously created template, and click Deploy Template.
15. Select Deploy to Devices, then click Select Sleds.
16. Select the sleds to receive the template, click Finish.
17. Click Boot to Network ISO in the left column, enter the details for the CIFS or NFS ISO share to be used to install ESXi, click Next.
18. Adjust iDRAC management IPs if necessary, click Next.
19. Deselect any unneeded Target Attributes, click Next.
20. Ensure Run Now is selected, click Finish.
21. Wait for configuration process to complete.
22. Connect to virtual console on each host, and wait for ESXi installer to finish loading from ISO share.
23. Install ESXi on workload cluster nodes.
24. Once ESXi install is complete, reboot each host.
25. After reboot, press F2 on each host to configure the ESXi installations.
26. In the configuration menu on each host:
    - In Configure Management Network, set:
      - IPv4 address, Subnet Mask, and Default Gateway
      - Primary DNS (and secondary DNS, if available)
      - Hostname (in FQDN format)
      - Press Esc twice to return to previous menus, then select confirm to restart the management network.
    - In Troubleshooting Options, Enable SSH and ESXi Shell, then press Esc twice to return to the main ESXi info screen.
27. On each host:
    a. Log into the ESXi web interface.
    b. Click Manage in the Navigator pane, then System→Time & Date.
    c. Click Edit NTP Settings.
    d. Select Use Network Time Protocol.
    e. Change the startup policy to Start and stop with host.
    f. Enter at least one NTP server IP address in the text box.
    g. Click Save.
    h. Click Services, select ntpd, and click Start.
    i. Ensure that TSM and TSM-SSH are also running.
    j. Click Networking in the Navigator pane.

k. Select the Port groups tab.
l. Select the VM Network port group, and Edit Settings.
m. Set the VLAN ID to the environment's Management VLAN.
n. SSH into each host, run the following commands to generate new self-signed certificates for the changed host names.

```
/sbin/generate-certificates
/etc/init.d/hostd restart && /etc/init.d/vpxa restart
```

28. In the VMware Cloud Foundation SDDC Manager Dashboard, select Commission Hosts.
29. Verify all parameters are met, then click Select All and Proceed.
30. Under the Add new option, enter the FQDN for the first new host.
31. Select VSAN Storage Type.
32. Select the destination Network Pool Name.
33. Enter the User Name and Password for the ESXi host.
34. Click Add.
35. Repeat steps 30-34 until all hosts are added.
36. Click the Select All checkbox in the Hosts Added field, and confirm the fingerprints.
37. Click Validate All.
38. Once validation is complete, click Next to review, then Finish when complete.
39. Click Workload Domains, then +Workload Domain.
40. In Storage Selection, Select vSAN, then click Begin.
41. Enter the Virtual Infrastructure Name and the Organization Name.
42. Select Enable vSphere Lifecycle Manager Baselines, then click Next.
43. Enter the Cluster name, click Next.
44. Enter the FQDN and IP for the new workload domain's vCenter Server VM.
45. Enter the root password for the VM, click Next.
46. Select to create a new NSX Manager cluster.
47. Enter the VLAN for the NSX host overlay network.
48. Enter the NSX Manager Virtual IP, FQDN, and IP addresses and FQDNs for the NSX nodes.
49. Enter the NSX Manager password, click Next.
50. Select the desired vSAN configuration based on the available storage, click Next.
51. Select the Workload hosts commissioned in step 38, click Next.
52. Select available licenses to apply to the Workload Domain, click Next.
53. Review the Object Names, and confirm they are correct. Click Next.
54. Review the configuration details page, and verify all settings are correct. Click Next.
55. Once process completes, click Finish.

This project was commissioned by Dell EMC.